

What is claimed is:

1. A method for managing persistent connections between data processing units of a computer system, wherein a first data processing unit is connected to a second data processing unit to send requests to the second data processing unit for processing, the method comprising the steps of:

monitoring a communication delay period for requests transferred from the first data processing unit to the second data processing unit;

comparing the monitored delay period with a threshold communication delay period to determine whether the monitored communication delay period indicates a predefined performance condition; and,

in response to determining that the monitored communication delay period indicates a predefined performance condition, adjusting the number of connections between the first and second data processing units.

2. The method according to claim 1, wherein the comparing step comprises determining whether the monitored delay period exceeds a first threshold delay period, and wherein the step of adjusting the number of connections is responsive to determining that the monitored delay period exceeds the first threshold to establish at least one additional connection.

3. The method according to claim 1, wherein the comparing step comprises determining whether the monitored delay period is less than a second threshold delay period, and wherein the step of adjusting the number of connections is responsive to determining that the monitored delay period is less than the second threshold to close at least one connection.

4. The method according to claim 3, wherein the comparing step also comprises determining whether the monitored delay period exceeds a first threshold delay period, and wherein the step of adjusting the number of connections is responsive to determining

that the monitored delay period exceeds the first threshold to establish at least one additional connection.

5. The method according to claim 1, wherein the monitored delay period is  
5 calculated by:

computing a difference between (1) a timestamp associated with the transfer of the request from the first data processing unit to the second data processing unit and (2) a timestamp associated with the receipt at the first data processing unit of a response to said request from the second data processing unit, and

10 subtracting (3) a time period measured as the time processing the request within the second data processing unit.

6. The method according to claim 5, wherein the monitored communication delay period is averaged for a set of requests processed during a predefined time period prior to  
15 said step of comparing the monitored communication delay period with a threshold communication delay period.

7. The method according to claim 2, further comprising the step of calibrating the system to determine a first threshold communication delay period above which the  
20 establishment of at least one additional connection is expected to reduce the communication delay period.

8. The method according to claim 7, wherein the step of calibrating the system to determine a first threshold comprises:

25 monitoring communication delay periods and corresponding request throughput information for different numbers of concurrent clients and different numbers of persistent connections;

determining a minimum number of concurrent clients at which a predefined percentage increase in throughput can be achieved by increasing the number of persistent  
30 connections by an integer value,  $a$ , between the first and second data processing units; identifying, with reference to the monitored communication delay periods and

corresponding request throughput information, a communication delay period corresponding to the determined minimum number of concurrent clients.

9. The method according to claim 8, wherein the step of calibrating the system is performed separately for each of CPU-intensive requests and data-intensive requests, and the step of calibrating comprising the additional step of selecting a minimum communication delay period from the identified communication delay period for CPU-intensive requests and the identified communication delay period for data-intensive requests.

10. The method according to claim 9, further comprising the step of using said calibration of the system to determine a second threshold communication delay period below which the closing of at least one connection is not expected to significantly increase the communication delay period.

11. The method according to claim 10, wherein the first and second threshold communication delay periods are computed as percentage differences from a selected minimum communication delay period.

12. The method according to claim 2 further comprising the step, performed prior to the step of establishing at least one additional connection, of checking whether the adjusted number of connections would exceed a maximum permitted number of connections, wherein the step of establishing the at least one additional connection is performed only if the adjusted number of connections would not exceed the maximum permitted number.

13. The method according to claim 12, further comprising:  
monitoring communication delay periods and corresponding request throughput information for different numbers of concurrent clients and different numbers of persistent connections; and  
determining the maximum permitted number of connections by: identifying a

maximum throughput, for different numbers of concurrent clients, from the monitored request throughput information; identifying a number of persistent connections corresponding to said identified maximum throughput for each respective number of concurrent clients; and selecting a minimum from said identified numbers of persistent connections.

14. The method according to claim 1, wherein the first data processing unit is a front-end network gateway node of a cluster-based data processing system, and the second data processing unit is a back-end processing node of the cluster-based data processing system, the method comprising:

the gateway node receiving requests from a client requestor via a network, passing received requests to respective ones of a set of back-end processing nodes, receiving responses from the respective back-end processing nodes, and forwarding received responses to the client requestor; and

the back-end processing node processing requests received from the gateway node to generate responses, and forwarding the responses to the gateway node.

15. The method according to claim 14, wherein the cluster-based data processing system comprises a plurality of back-end processing nodes, and wherein the method is responsive to monitored communication delays for the plurality of back-end processing nodes to modify the number of persistent connections consistently for the plurality of back-end processing nodes.

16. A data processing system comprising:

a first data processing unit for receiving requests from a client requestor and passing received requests to a second data processing unit, and for receiving responses from the second data processing unit and forwarding received responses to the client requestor; and

a second data processing unit for processing requests received from the first data processing unit to generate responses, and for forwarding the responses to the first data processing unit; and

a connection manager for managing the number of connections between the first and second data processing units, the connection manager being responsive to monitored communication delays between the first and second data processing units indicating a predefined performance condition to modify the number of persistent connections between the first and second data processing units.

17. The data processing system according to claim 16, wherein

the second data processing unit comprises a monitor for monitoring and recording a processing time period corresponding to the time taken to process a request at the second data processing unit, and

the first data processing unit comprises a monitor for recording a first timestamp and a second timestamp, the first timestamp generated when sending a request from the first data processing unit to the second data processing unit and the second timestamp generated when a response to said request is received by the first data processing unit from the second data processing unit, for determining a difference between the first timestamp and the second timestamp, and for subtracting said processing time period from said difference to compute a communication delay period.

18. A data processing system comprising:

a network subsystem;

a gateway node for receiving requests from a client requestor and passing received requests to the network subsystem for delivery to one of a set of back-end processing nodes, and for receiving responses from the back-end processing nodes via the network subsystem and forwarding received responses to the client requestor; and

a set of back-end processing nodes for processing requests received from the gateway node via the network subsystem to generate responses, and for forwarding the responses to the gateway node via the network subsystem;

a connection manager for managing the number of connections between the gateway node and each of the back-end processing nodes, the connection manager being responsive to monitored communication delays between the gateway node and the back-end processing nodes indicating a predefined performance condition to modify the

number of persistent connections between the gateway node and at least a first one of the back-end processing nodes.

19. The data processing system according to claim 18, wherein the connection  
5 manager is responsive to monitored communication delays for the set of back-end processing nodes to modify the number of persistent connections consistently for all back-end processing nodes of the set.

20. The data processing system according to claim 18, wherein  
10 each of the set of back-end processing nodes includes a monitor for monitoring and recording a processing time period corresponding to the time taken to process a request at the respective back-end processing node, and  
the gateway node includes a monitor for recording a first timestamp and a second timestamp, the first timestamp generated when sending a request from the gateway node  
15 to a respective one of said back-end processing nodes and the second timestamp generated when a response to said request is received by the gateway node from the respective back-end processing node, for determining a difference between the first timestamp and the second timestamp, and for subtracting said processing time period from said difference to compute a communication delay period.

20

21. The data processing system according to claim 18, wherein the connection  
manager includes means for comparing monitored communication delays with a first  
threshold communication delay period, and wherein the connection manager is  
responsive to said comparing step determining that the monitored communication delays  
25 exceed the first threshold to establish at least one additional persistent connection between the gateway node and each of the back-end processing nodes.

22. The data processing system according to claim 18, wherein the connection  
manager includes means for comparing monitored communication delays with a second  
30 threshold communication delay period, and wherein the connection manager is responsive to said

comparing step determining that the monitored communication delays are less than the second threshold delay period to close at least one connection between the gateway node and each of the back-end processing nodes.

5 23. The data processing system according to claim 18, wherein the set of back-end processing nodes are nodes of a clustered Web server.

24. A method for managing persistent connections between a gateway node and each of a set of back-end processing nodes arranged in a cluster within a data processing  
10 system, the method comprising the steps of:

monitoring a communication delay period for requests transferred from the gateway node to the back-end processing nodes;

comparing the monitored communication delay period with a threshold communication delay period to determine whether the monitored communication delay  
15 period indicates a predefined performance condition; and,

in response to determining that the monitored communication delay period indicates a predefined performance condition, adjusting the number of connections between the gateway node and at least one of the set of back-end nodes.

20 25. The method according to claim 24, wherein the comparing step comprises determining whether the monitored delay period exceeds a first threshold delay period, and wherein the step of adjusting the number of connections is performed to establish at least one additional connection in response to determining that the monitored delay period exceeds the first threshold.

25 26. The method according to claim 24, wherein the comparing step comprises determining whether the monitored delay period is less than a second threshold delay period, and wherein the step of adjusting the number of connections is performed to close  
30 than the second threshold.

27. The method according to claim 24, wherein the monitored communication delay period is calculated by:

computing a difference between (1) a timestamp associated with the transfer of the request from the gateway node to a respective back-end data processing node and (2)  
5 a timestamp associated with the receipt at the gateway node of a response to said request from the respective back-end processing node, and

subtracting (3) a time period measured as the time processing the request within the back-end processing node.

10 28. The method according to claim 27, wherein an average monitored communication delay period is calculated for a set of requests processed during a predefined time period prior to said step of comparing the monitored delay period with a threshold delay period, and wherein comparing the monitored delay period comprises comparing the calculated average monitored delay period with the threshold delay period.

15

29. The method according to claim 25, further comprising the step of calibrating the system to determine a first threshold communication delay period above which the establishment of at least one additional connection is expected to reduce the communication delay period.

20

30. The method according to claim 25 further comprising the step, performed prior to the step of establishing at least one additional connection, of checking whether the adjusted number of connections would exceed a maximum permitted number of connections, wherein the step of establishing the at least one additional connection is  
25 performed only if the adjusted number of connections would not exceed the maximum permitted number.

31. The method according to claim 30, further comprising:  
monitoring communication delay periods and corresponding request throughput  
30 information for different numbers of concurrent clients and different numbers of persistent connections; and



determining the maximum permitted number of connections by: identifying a maximum throughput, for different numbers of concurrent clients, from the monitored request throughput information; identifying a number of persistent connections corresponding to said identified maximum throughput for each respective number of concurrent clients; and selecting a minimum from said identified numbers of persistent connections.

32. A computer program product, comprising program code recorded on a recording medium for controlling operations on a data processing system on which the program code executes, the program code comprising a connection manager for managing the number of connections between a first data processing unit and a second data processing unit of the system by:

monitoring a communication delay period for requests transferred from the first data processing unit to the second data processing unit,

comparing the monitored communication delay period with a threshold communication delay period to determine whether the monitored communication delay period indicates a predefined performance condition; and,

in response to determining that the monitored delay period indicates a predefined performance condition, adjusting the number of connections between the first and second data processing units.

33. A data processing system comprising:

a gateway node for receiving requests from a client requestor and passing received requests to one of a set of back-end processing nodes, and for receiving responses from the back-end processing nodes and forwarding received responses to the client requestor; and

a set of back-end processing nodes arranged in a cluster, for processing requests received from the gateway node to generate responses, and for forwarding the responses to the gateway node;

a connection manager for managing the number of connections between the gateway node and each of the back-end processing nodes, the connection manager being

responsive to monitored communication delays between the gateway node and the back-end processing nodes indicating a predefined performance condition to modify the number of persistent connections between the gateway node and at least a first one of the back-end processing nodes.